


Sun Grid Engine @ Synopsys

Joe Fu
IT Manager, Office of CIO



Agenda

- Company profile
- Compute environment overview
- SGE management
- SGE key benefits
- Challenges

Introduction to Synopsys

Synopsys (Nasdaq: snps) is a worldwide leader in electronic design automation (EDA) software and IP for semiconductor design. Synopsys addresses the key challenges designers and manufacturers face today, including power management , accelerated time to yield and system-to-silicon verification.

<http://www.synopsys.com/>

What does that mean ?

- We develop software using common dev tools like C/C++ compilers, debuggers, perl etc
- We support several Unix/Linux/Windows platforms
- The products need to pass comprehensive regression Q/A before release
- Some EDA tools are capable of partitioning the problem into thousands smaller ones in parallel/distributed computing environment.
- We are a global company and requires a model that allows optimum usage of all resources

Overcome Early Challenges

- **Geo – Social Political Challenges**

- Taking Machine away from smaller offices to relocate into Central clusters
- Converting desktop computing to server class computing
- Lack of long term trending data for resource planning
- Resistance to move away from direct login work model

- **Technical Challenges**

- Lack of tools/ utilities to help IT/customers effectively manage/utilize compute resources
- Performance and scalability issues of scheduler software

Compute Environment Evolution


	2000	2002	2004	2006	2008	2010
cores	1000	4000	8000	10k	20k	40k
Regression jobs	10k	50k	200k	500k	1.5 M	3M
Regression method	manual	cron based	1 / 2 cron & half grid	grid based nightly	24x7 grid based	24x7 grid based

- Over the past few years
 - CPU Performance has continuously doubled.
 - CPU/storage becomes much more cost-effective.
 - Computing hardware acquisition has exploded
 - Computing infrastructure becomes enormous and more complex


Facts and Figures

- Over 90% computing cores in scheduler based clusters with 70% overall average CPU utilization
- All jobs; Interactive, Regression, Benchmark, Build and Production are done via compute clusters
- Over 15k cores (CPUs) in SGE compute clusters worldwide with various configurations
- Support over 3000 R & D engineers, application consultants and professional services
- Over 2500 cores in the largest farm (SGE) and growing constantly
- Daily throughput ranges in the order of several million on the compute clusters

Compute Cluster Profile



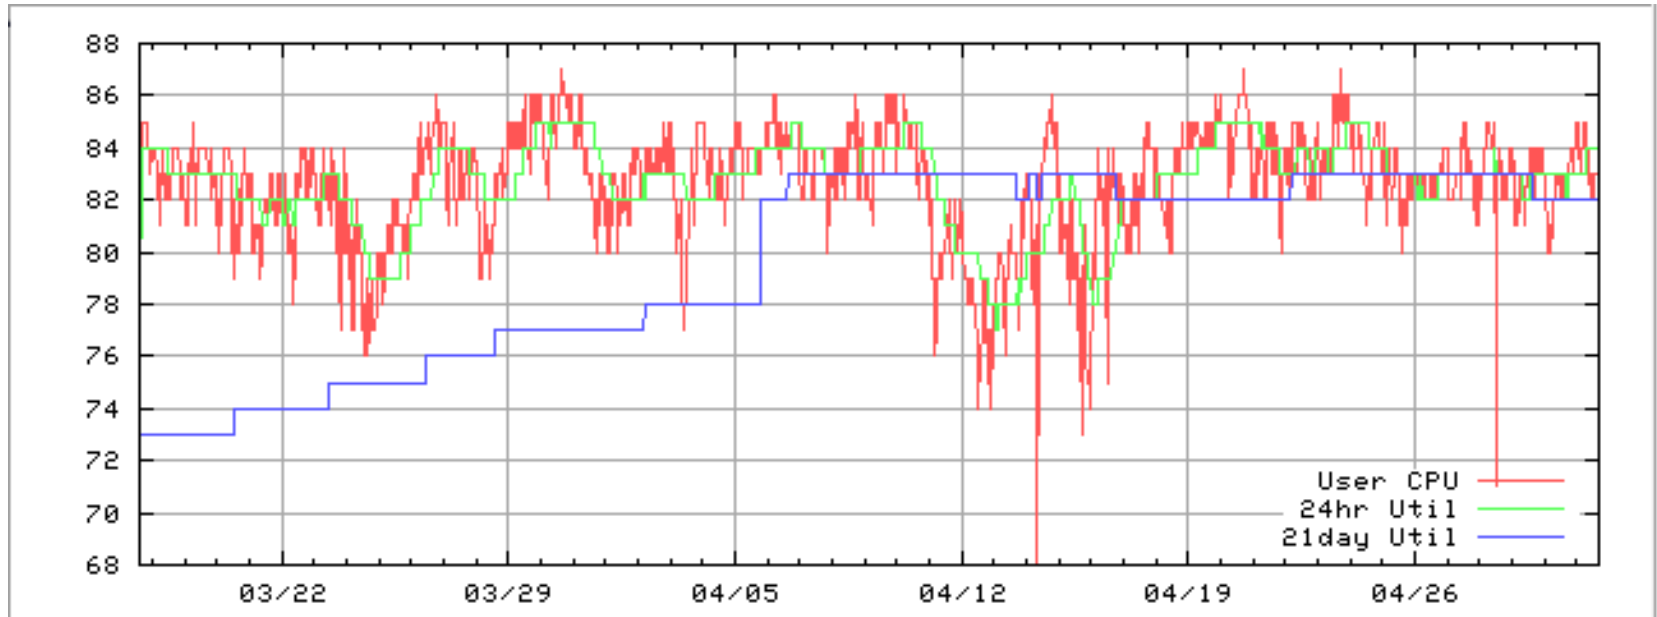
General Purpose Compute



High Performance Cluster

Compute Cluster Utilization

Over 2500 Cores @ Approximate CPU utilization of 80%




Achieve Optimum Utilization


- Compute farm self-healing framework to avoid black-hole machines and suspend bad jobs
- Auto healthy/unhealthy system check by load sensor to dispatch jobs only to healthy machines to prevent job failures
- Memory alarm implementation and memory requirement enforcement to prevent resource contention
- Dynamic resource quota control to guarantee regression throughput for multiple products
- ‘rush’ access to certain resources by using pre-empt feature or unix priority/nice feature so jobs can coexist

Exception to the norm..

- large 128GB machine (burst of usage)



- Light interactive Machine (allow idle jobs)



Business Rules Monitoring

- Service Level Agreement (SLA)
 - Compute clusters required to guarantee daily job throughput so products can be released on time
 - Compute clusters required to sustain >95% overall slots utilization
- Hurdles to SLA
 - Users steal resource by direct login
 - Users allow computing slots to sit idle
 - Long/run-away jobs wastes slots
 - Swapping/thrashing compute resources

BRM ENSURES SLA

SGE Key Benefits

- Scalability : one of the biggest farms has tripled the number of cores over two years and we plan to double the size the farm in next year. Good performance and responsiveness given the size of the farm and the number of the jobs
- Flexibility: the individual queue instance and multi-fairshare for different projects, user access control, resource quota restriction allow to provide different farm methodologies to meet R & D needs
- Easy Administration: 90% of our IT administration work are automated because SGE provides powerful 'qconf' configuration utility
- Easy to use: minimum training will get users start working in SGE and the convenience to add any submit host on the same network
- Low cost: relatively inexpensive compared to other more commercialized software such as LSF

Challenges in Compute Clusters

- Multi-thread/multi-core jobs become popular
- More comprehensive usage/rules monitoring and enforcement system to provide flexibility as well as efficiency of resource usage
- Multi-layer and multi-option to manage large resources as EDA jobs become larger and larger
- Automated Provisioning to support multiple on demand operation system configurations

SYNOPSYS®

Predictable Success